Amendments to the Specification

Please replace the paragraph at page 13, lines 4 through 6 with the following amended paragraph:

Optionally, there is a load balancing function [[11]] 16 in front of the host 12 processors, which directs individual transactions to specific host or hosts 12 so as to evenly distribute workload.

Please replace the paragraph at page 17, line 17 with the following amended paragraph:

[[B.]] C. DETAILED DESCRIPTION OF PSDP ARCHITECTURE

Please add the following new paragraphs at page 17, line 17 (before B. DETAILED DESCRIPTION OF PSDP ARCHITECTURE):

B. JPU SOFTWARE COMPONENTS

Fig. 3 is a diagram of the software components of a JPU 22.

Communications Layer 300

- Provides internal communication among nodes
- Includes Job Listener 301 to await requests
- Includes Network Poster 302 to send data when buffer filled, job completed, or at Host request

JPU Dispatch/Scheduler 304

- Receives plan through Communications Layer 300
- Queues Plan
- Schedules/dispatches jobs according to their priority, "fairness" to date, expected resource requirements, and available resources

JPU Transaction Manager 306

- Processes changes in transaction state to begin a transaction, pre-commit a transaction, commit a transaction, or abort a transaction
- Handles processing of dependencies among transactions as flagged by the lock manager; broadcasts information about these dependencies to relevant host(s); initiates deadlock checks

JPU Lock Manager 308

- Controls concurrent access to data
- Interfaces with EventTask 310 before a query is executed and for each result set returned from a scan
- Provides support for arithmetic locking

JPU Recovery Manager 312

- Maintains a Journal to track transaction status on the JPU 22, using the Storage Manager API
- Performs transaction recovery when requested by JPU Transaction Manager

JPU Mirror Manager 314

- Mirror Sender receives copies of record updates from Storage Manager 320 and transmits these to the mirror for this JPU when an updating transaction commits
- Mirror Receiver receives record updates, buffers these in memory, and flushes out to disk through the Storage Manager when the Mirror Receiver buffer is full
- Transmits all data to a spare system during failover processing

Storage Manager 320

- Stores and manages information on disk in optimal fashion
- Has an API that supports storage and retrieval of records (or tuple sets)
- Supports error checking to insure that the data conforms to the indicated table and the indicated table "owns" the physical space to which the data is being written

- Supports creation and deletion of tables, views, and indices
- Handles record inserts and deletes
- Supports ETL and mass loading of existing user data among various JPUs
- Provides storage support for commit/rollback
- Provides support for Precise Indexes
- Provides mirroring support for failover
- Optimizes sort operations and utilizes smart hash algorithm for data distribution/striping
- Provides support for compression and smart storage optimization
- Controls disk I/O

JPU Resource Scheduler 322

- Schedules jobs to run on the PSDP 28; communicates with JPU/PSDP Scheduler 324 to queue up PSDP requests to retrieve required data
- Optimizes the queue to keep the PSDP/disk as busy as possible, with requests from multiple queries intermixed in the queue based on disk characteristics and location of data on the disk
- Takes into account the needs of any data loading for new tables being created and transformed to internal data format (i.e., to optimize the loading process)
- Supports heuristic-based scheduling, ensuring that jobs are scheduled on a priority basis, but also ensuring that all jobs do get serviced (e.g., raising a job in priority if it has not been run in a certain interval of time)
- Supports synchronous/piggy-backed scans, combining similar requests to optimize PSDP processing
- Manages memory buffers/memory allocation on JPU; allocates memory to Execution Plans based on expected needs and hints received from Plan Optimizer
- JPU Paging (if required)

PSDP Prep 330

· Defines the instructions that will be given to the PSDP 28 in order to process a request (instructions tell the PSDP 28 what to do with each field being read from the disk)

· Identifies what filtering, transformation, projection, and aggregation operations are to by run by the PSDP 28

EventTask 310

· Executes the portion of the Execution Plan that could not be handled by the PSDP but that does not have to be handled at the Host level

· Handles sorts, joins, transformations, and aggregations that could not be done as data stream through the PSDP 28

· Maintains a memory buffer of result set records and returns these to Host through the Comm Layer when buffer filled, job completed, or at Host request

JPU Diags 332

· Runs diagnostics on JPU as required/requested

JPU Boot/Init 334

· Executes image burned into flash memory at boot time to bootstrap the JPU, run diagnostics, register the JPU with the primary Host server, and download new image from Host to run

· Loads and transfers control to the image downloaded from the primary Host server to load the JPU application code, the operating system, the network stack, and disk driver code

Backup/Recovery 336

· Supports JPU side of Backup/Recovery process

· Interfaces with Transaction Manager and JPU Storage Manager

DBA Lite 338

· Provides automatic and dynamic disk and Storage Manager support

- Supports dynamic index creation, defragging, index garbage collection, timers, agents

JPU/PSDP Scheduler 324

- Schedules jobs to run on the PSDP; queues up PSDP requests to retrieve required data